

US-PAT-NO: 6421767

DOCUMENT-IDENTIFIER: US 6421767 B1

**TITLE: Method and apparatus for managing a storage system
using snapshot copy
operations with snap groups**

DATE-ISSUED: July 16, 2002

**US-CL-CURRENT: 711/162; 707/204 ; 711/111 ; 711/112 ; 711/114 ;
711/152
; 711/161**

APPL-NO: 09/ 511886

DATE FILED: February 23, 2000

----- KWIC -----

Brief Summary Text - BSTX:

A data file snapshot copy is an improvement over this type of copy process.

This snapshot copy process includes a dynamically mapped virtual data storage

subsystem. This subsystem stores data files received from a processor in

back-end data storage devices by mapping the processor assigned data file

identifier to a logical address that identifies the physical storage location

of the data. This dynamically mapped virtual data storage subsystem performs a copy of a data file by creating a duplicate data file pointer to a data file identifier in a mapping table to reference the original data file. In this dynamically mapped virtual data storage subsystem, the data files are referred to as a collection of "virtual tracks" and each data file is identified by unique virtual track addresses (VTAs). The use of a mapping table provides the opportunity to replace the process of copying the entirety of a data file in the data storage devices with a process that manipulates the contents of the mapping table. A data file appears to have been copied if the name used to identify the original data file and the name used to identify the copy data file are both mapped to the same physical data storage location.

Detailed Description Text - DETX:

Further, the use of snap groups restricts which virtual volumes are allowed to be paired for a snapshot copy operation. According to the present invention, when selecting a source virtual volume for a snapshot copy operation, the target virtual volume must be a virtual volume within the same snap group rather than any virtual volume in the storage subsystem.

Detailed Description Text - DETX:

These virtual track table pages each contain an entry for each virtual track.

Also located within each virtual track table page is data, which defines the

logical address of a copy of the virtual track table page comprising a virtual

track table page instance, which has been written on back-end data storage

devices during the snapshot copy or write operation. These back-end storage

devices may be, for example, storage devices 202 in storage subsystem 200 in

FIG. 2. This logical address identifies the physical storage location in the

back-end data storage devices that contains the most recently written instance

of the present virtual track table page.

Detailed Description Text - DETX:

A track number table page address points to a track number table page, which

contains a predetermined number (for example: 8192) of byte segments of memory.

These track number table pages each contain an entry for each virtual track

within a 1024 track number boundary. As with the virtual track table, the

physical storage for these virtual track tables may be within a cache memory,

within a controller, within data storage devices, or a combination of these

locations as a matter of engineering choice.

Detailed Description Text - DETX:

The process begins by performing a snapshot validation process (step 1000).

This validation process in step 1000 refers to checks made by the subsystem to

verify that the snapshot copy request follows the rules required to perform a

successful snapshot copy operation. For a general snapshot copy operation, a

number of checks are typically performed. These checks include, for example:

(1) ensuring that the source volume and the target volume are configured, (2)

determining whether the extents for the source volume and the target volume are

within the limits of the volume, (3) whether the user is able to issue a "snap

from" command to the appropriate extent on the source volume,

(4) whether the

user is able to issue a "snap to" command to the appropriate extent on the

target volume, (5) insuring that the extents specified for the source volume

and the target volume have the same number of tracks, and (6)

ensuring that no

part of both extents are involved with a different snapshot copy operation

still in progress.

Detailed Description Text - DETX:

Basically, reference count regeneration needs to scan all the

virtual tracks in the subsystem and maintain a count of which physical tracks are referenced by each virtual track. With the current implementation without snap groups, no mechanism is present to predict or limit which physical track a particular virtual track will reference. As a result, either counters must be maintained for all possible virtual tracks on a single scan of the virtual tracks or the virtual tracks have to be scanned more than once depending on the number of counters that are maintained. Maintaining a counter for each physical track consumes large amounts of memory and can be a limiting factor on the number of tracks supported. Scanning the virtual tracks involves reading virtual track information from physical storage and is a very slow process. Multiple scans of the virtual tracks creates an unacceptable performance degradation for this operation.

Detailed Description Text - DETX:

With the snap group implementation of the present invention, the number of counters to maintain is controlled by the number of tracks in the snap group. Each counter is associated with a physical track in these examples. The scan of the virtual tracks within the snap group cannot contain a reference to a

physical track outside of the snap group. The number of physical tracks cannot be more than the number of virtual tracks in a snap group in the present invention. Thus, the number of virtual tracks forms the upper limit for the number of possible physical tracks that may be associated with a counter. As a result, the number of counters is limited and the scan of the entire virtual track information is performed once with the scan broken into "n" pieces where "n" is the number of snap groups.